




# Six reasons to upgrade your data management

How to become more data driven by  
managing growth and data proliferation  
while cutting cost






# Contents



This ebook presents  
six reasons why you  
should upgrade your  
database, including  
opinions from industry  
analysts and real-world  
customer experiences.

[Get started >](#)





# Data isn't the same as it used to be.

## Your database shouldn't either.

The days of data being narrowly defined as highly structured information from a few specific sources is long gone. Replacing that notion is the reality of a wide variety of data types coming from multiple sources, internal and external, to an organization. All of it is in service of providing everyone from IT, to line-of-business (LOB) employees, to C-level executives with insights that can have an immediate and transformative impact.

Effectively using and managing information is critical to pursuing new business opportunities, attracting and retaining customers, and streamlining operations. To do so, organizations must accommodate a rapidly increasing volume, variety and velocity of data while extracting actionable insights faster than ever. However, these needs create an array of workload challenges and increasing demands on underlying IT infrastructure and database systems that are often not up to the task.

The question is, how will you solve for these challenges? Will you allocate more staff to keep up with the patches, add-ons and continual tuning required by existing systems? Simply ignore the potential insights that lay in this wealth of new data? Or will you find new database solutions that enable you to innovate and move your data infrastructure and business forward?



## Reason 1

# Lower your total cost of ownership

It is more important than ever that your IT team uses its budget and staff efficiently. You need a database that cost-effectively provides industry-leading data management capabilities while also meeting SLAs with the minimum possible headcount.

Automating administrative tasks helps with both those goals. Look for a database solution that can automate a range of administrative tasks, such as setup and deployment, workload management, resource utilization and storage management, as well as maintenance, upgrades, and capacity expansion.

Data compression and multi-temperature data management technologies offer another cost-saving tactic because they efficiently reduce storage requirements that often consume a large percentage of the IT budget.

Additionally, cluster topology transparency—not requiring applications to be aware of the underlying cluster and database deployment—helps accelerate coding and testing, and can make application developers more productive. This allows you to switch databases to save money and accelerate processes, without having to modify current applications.

### Compare the costs and benefits of two database leaders for online transactional processing (OLTP) deployments

Learn more about database costs in this report from analyst group Quark + Lepton. The report focuses on high-performance analytics and transaction processing, including comparing cost of ownership between Oracle and IBM.

[Read the report](#)




32%

Find out which database averaged **32% less** on three-year costs of transaction processing environments.



“By migrating our core  
insurance administration system  
[...] we have cut licensing  
and database administration  
costs by 80%.”

—Jean-Yves Pignier, CIO, Groupe FranceMutuelle





## Reason 2

# A platform for rapid reporting and analytics

**In order to compete today**, organizations need to get insights from their data, fast. Data-driven insights can drive organization-wide innovation, uncovering opportunities for new products or markets, empowering salespeople to have more meaningful discussions, and identifying internal processes that can be improved. For that reason, many organizations have turned to databases with in-memory columnar technology, which can speed reporting and analytics through data compression and data skipping.

**In-memory columnar technology** is particularly important for data scientists and business analysts—key players in delivering organizational insight and innovation. They need an environment where they can rapidly work with data to arrive at meaningful answers. Active compression and data skipping help in this regard. The queries a data scientist or business analyst might run can be returned faster when the data doesn't need to be decompressed before use and irrelevant data isn't used.

**Your LOB users also need quick answers** to their questions. Fortunately, the same technology that helps the data scientists and business analysts also helps deliver reports more quickly to this critical group, which includes executives, mid-level managers and customer-facing personnel.



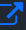


# Is your reporting built for the Cognitive Era?



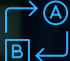

The key to accelerating insight is providing a range of data processing techniques designed to deliver fast answers in an easy-to-use format.



As organizations prepare for the cognitive future, it's vital to have a solid database foundation that rapidly delivers insights. In particular, data scientists need access to data more quickly to test theories and deliver the fresh perspectives.


The paper *Is your database ready for the Cognitive Era?* explores in detail the steps that businesses must take to remain competitive. It also discusses the role that in-memory columnar technology plays in delivering insights quickly.

[Read the full paper here](#) 


## Check if your database has the following:

-  Load-and-go capabilities
-  Actionable data compression and data skipping
-  Parallel processing techniques
-  High performance for complex queries out of the box

-  Multiplatform flexibility
-  No need for specialized hardware configurations

-  Exploitation of single instruction, multiple data (SIMD) processing





"We have shrunk data loading times by 50% on average—with some data loads running 12x faster than before. What's more, we have accelerated queries by 20% on average."

—Micha Altmeyer, Team Leader  
SAP Application Management Services, Helvetia





### Reason 3

# Increase scalability and availability

Keeping your business operating requires data systems that are up and running all the time and strong enough to handle the increasing amount of transactional data generated by your enterprise. Your customers and partners expect continuous uptime and instantaneous responses. In this social media age, poor performance can quickly turn into a public relations nightmare.

While scaling up to accommodate the proliferation of data can be accomplished through different styles of clustering, the most appropriate style for transactional data is shared storage clustering. Shared storage clustering allows work to be distributed among different nodes that all rely upon the same data source. This process provides an opportunity to better deliver high availability.



**When considering upgrading your database,** look for two key shared storage clustering capabilities that will enable you to seamlessly handle unplanned or external disruptions with redundancy, high availability and disaster recovery:



# 1

The implementation you choose should be able to deliver automatic workload balancing across the nodes in your architecture.

# 2

You need a data system with built-in automatic failover to use each node in the shared storage cluster as a failover node for another.


**Knorr-Bremse turns to a new database to help with increasing data volumes**

For Knorr-Bremse, increasing its global presence through acquisitions and joint ventures meant a large increase in data volumes. Running a new database with enterprise strength unlocked its ability to scale.

[Read the Knorr-Bremse case study to learn how.](#)

“We can dynamically scale resources up and down as needed,  
without any impact on business users.”

— **Marc Moteja, Director**, IT Technology & Infrastructure, Knorr-Bremse





## Reason 4

# Support new data types and sources

Multiple data types and sources are now more prominent. Consider the dramatic growth of semi-structured and unstructured data generated from sources, such as social media, sensors in machinery connected to the Internet of Things (IoT), and audiovisual content repositories.

As the value of this information becomes apparent, businesses must bring the data that's outside their organization together with what already exists. However, this process poses a new set of deployment and development challenges for the database infrastructure, which needs to provide the flexibility to serve, capture, store and process a wide variety of information types from different sources. You need a data store that can serve data when and how it's needed, and capture exactly what's happening at any moment in time.



**Audiovisual content** provides the opportunity for organizations to tap into knowledge that might otherwise be difficult to find or review, highlighting new areas of potential profit.



**Social media data** can inform a 360° view of the customer, helping target marketing and drive increased sales.



**IoT sensor data** can provide alerts when machinery or vehicles are at risk of malfunction, encouraging preventative maintenance as opposed to emergency repairs.

**Your database must integrate well with open source tools, such as Apache Hadoop. Hadoop is often the primary repository for unstructured information due to its ability to avoid assigning a schema when storing data.**

When considering a database upgrade, make sure your technology partner has a close relationship with top Hadoop distributors and the open source communities. Partners with strong connections are more prepared to help you seamlessly connect with the unstructured data in Hadoop. Make sure your technology partner can describe the process by which data that's captured directly within Hadoop without a schema is queried and integrated with data stored in existing data warehouses or databases. Similarly, ask how "cold" data can be offloaded into Hadoop from a data warehouse without delays or excess effort.

## Can your database support all types of data?

- Does the database vendor have a strong relationship with the open source community and Hadoop providers?
- Can the database integrate well with open source technologies, such as Hadoop?
- How well does the database interact with other components of your data ecosystem, such as the data warehouse?



## Reason 5

# Increase the flexibility of your deployment

Depending on your unique needs, consider which deployments best meet your requirements: on-premises private, hosted or managed cloud.

See if your current database is ideally suited to your needs and work with a trusted technology partner to select any new deployments. Each deployment has its strengths and drawbacks, which is why you need the flexibility to select the one that best meets your goals and addresses your organization's unique business environment.



### On-premises

On-premises deployments provide a high level of control over optimization, data access and visibility, but require large investments in hardware and IT talent. They're most often used when security is paramount or government regulations mandate such deployments.



### Private cloud

Private cloud deployments allow you to take advantage of the security of your own firewall while providing flexibility traditionally associated with the public cloud. The ability to use container technology is a considerable benefit for this type of deployment. Compared to virtualized machines, they are easier to deploy quickly and often require less server space to operate.



### Hosted cloud


Hosted cloud environments don't provide as much control over bandwidth or optimization, but still allow you to manage the system to some degree. A hosted cloud is a good option for disaster recovery projects, as you can customize them to reflect the on-premises database more exactly without the added expense of provisioning more on-premises servers.



### Managed cloud


Managed cloud deployments provide the lowest amount of direct control over aspects like bandwidth and monitoring, but have heightened flexibility. Rapid elasticity, known as cloud bursting, separating compute and storage needs, and paying only for what you use mean that managed cloud deployments are great options for backups and development, or testing environments.





# Choosing the right database for your hybrid environment.

Most companies now deploy their database in a variety of ways. Yet, choosing a database that excels in that type of hybrid environment requires serious consideration. The paper *Bridging to the Cloud* discusses a full list of database requirements for moving to the cloud, such as support for mixed workloads and all data types.

Information about one database's ability to act as your bridge to the cloud is also **provided**. Read the paper, *Bridging to the Cloud* [here](#). 



## Reason 6

# Secure your database as much as possible

Don't look at the security of the database in a vacuum. The enterprise-strong performance of the database should be maintained as much as possible despite rigorous security measures.

Ideally, databases should demonstrate that their encryption overhead can reach single digits under optimal conditions.

With the frequent reports of data breaches in the news and the damaging effects that they can have on the public's goodwill toward a company, it's no surprise that database security is receiving extra scrutiny. As one of your company's most valuable assets, protecting your data with strong, enterprise-ready security is a top priority.

There are many ways to help safeguard your data, but one of the most important is encryption for data in transit and encryption for data at rest. Your database should provide both.

### Data in transit

Data in transit between a client application and a database server isn't the only line of communication that should be encrypted using secure sockets layer (SSL) technology. Those taking advantage of high availability disaster recovery (HADR) will also want to use SSL when logs are shipped from the primary server to standby servers.

### Data at rest

Native encryption of data at rest, including backup images, is vital for a database. All sensitive data should be encrypted. Look for databases that have support for Key Management Interoperability Protocol (KMIP) 1.1 compliant centralized key managers. Also, look for native encryption capabilities that can exploit third-party hardware security modules (HSMs) to safely store master keys.



# How to make upgrading easier



The way organizations depend on data is deepening with every day. This resource is growing rapidly, with more types and sources than ever before, presenting opportunities and challenges. That's why it's crucial to have a database and infrastructure that address these needs and grow with them.

If your database doesn't measure up, now is the time to think about upgrading. Reach out to a technology partner that can not only provide a database ready for the stringent demands of today, but also offer guidance and support for any questions you may have in the future. By working with a partner that has a wide range of data management options and a strong history of integration, organizations can build a foundation based on forward looking flexibility.

A good place to start your search is with IBM. As a leader in database technology, IBM can offer the assistance you need and supply a database that's prepared for current and future data challenges in IBM® Db2® database servers.

**To learn more about IBM Db2 database servers and how IBM makes it easy to upgrade, contact your IBM representative or IBM Business Partner, or visit [ibm.com/db2](https://ibm.com/db2).**



© Copyright IBM Corporation 2018

IBM Corporation  
Route 100  
Somers, NY 10589

Produced in the United States of America  
March 2018

IBM, the IBM logo, ibm.com, Cognitive Era, and Db2 are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.